

VACET Visualization and Analysis at ORNL (2007 update)

Sean Ahern

www.vacet.org

ORNL VACET Team



Sean Ahern
 Dave Pugmire (new)
 George Ostrouchov
 Jeremy Meredith
 Jamison Daniel
 Markus Glatter (UTK student)

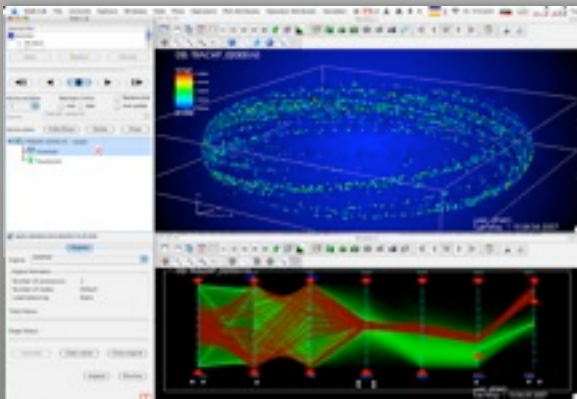
ORNL VACET Compute Resources

- 2007: 100,000 hours on Cray XT3/4 (Jaguar):
 - #2 on Top500 list at 119 TF
 - 23,016 Opteron cores
 - Many SciDAC customers
 - *Ideal for in-situ analysis work*
- Existing vis system (Hawk):
 - 44 node vis cluster, 88 processors, 88 GB RAM
- About to site new vis/analysis cluster (Lens):
 - 32 nodes, 512 Opteron cores, 2 TB of RAM, 64 GPUs
 - Connected to large Lustre store (~1.6 PB)

Some notable projects in the last year

High dimensional filtering

- Information visualization technique (a.k.a. parallel coordinates)
- Consists of three linked capabilities
 - Parallel coordinates plot with summary view for large data
 - Restriction tool
 - Multivariate threshold operator
- Implemented in VisIt
- Used to filter multivariate data from particle-in-cell fusion code
- Demonstrated in parallel on ~60 million zone combustion data



Efficient Boolean Range Queries

- Harnessed VisIt's "contracts" system to optimize multivariate data thresholding
- Connected to IUSV technology for scalable data servers
 - "Scalable Data Servers for Large Multivariate Volume Visualization", Markus Glatter, Colin Mollenhour, Jian Huang and Jinzhu Gao, IEEE Transactions on Visualization and Computer Graphics, Vol. 12, No. 5, pp. 1291-1299, 2006.
- Puts groundwork in place for connecting to other query systems (FastBit)
- *Will be demoed at Supercomputing in 2 weeks*




Representative Particle Field Sampling

- Statistical methods to reduce I/O for extremely large datasets
- Preserves nature of data, including trends and extrema
- R infrastructure:
 - Uncertainty quantification
 - Cluster algorithms
 - Extreme value theory
 - Data parallel implementation (R-MPI)

VisIt work (lower level)

- Assisted in move from LLNL to NERSC code repository
 - Have successfully transitioned entirely from ClearCase at LLNL to Subversion at NERSC
- Database options
 - Provides ability to create GUIs on the fly for custom data loading and parameterized I/O
 - More work needed here
- Automatic SSH tunneling
 - Needed for remote deployment in firewalled environments
- Debugging APDEC/Chombo data from within gdb using VisIt



Upcoming work

- Deployment of all major VACET tools on new systems (VisIt, ViSUS, Manta, SCIRun)
- New fusion support in VisIt
 - Streamline improvements
 - Poincaré analysis
 - Temporal analysis
- Deployment of CDAT/ViSUS workflow for Drake (climate)
- In-situ analysis for Chen (combustion) in VisIt
- Move entire VisIt development process to GForge (wiki, testing, mailing lists, bug tracking, ...)

